CROWD ANALYTICS VIA ONE SHOT LEARNING

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH & DEVELOPMENT

This invention was made with Government support under contract number 2013-11-CX-K010 awarded by the National Institute of Justice of the United States Department of Justice. The Government has certain rights in the invention.

BACKGROUND

The subject matter disclosed herein relates to crowd analytics via one shot learning.

Understanding individual or crowd level behavior is an important field of study and may lead to behavior detection. Detecting a person's behavior may enable crime reduction and/or enhanced security in various locations where crowds typically gather, such as airports, train stations, sporting arenas, movie theaters, and the like. It is now generally recognized that improved techniques to detect a behavior prior to a person carrying out that behavior is desirable.

BRIEF DESCRIPTION

Certain embodiments commensurate in scope with the originally claimed subject matter are summarized below. These embodiments are not intended to limit the scope of the claimed subject matter, but rather these embodiments are intended only to provide a brief summary of possible forms of the present disclosure. Indeed, the disclosed techniques may encompass a variety of forms that may be similar to or different from the embodiments set forth below.

In one embodiment, a method is provided that includes the steps of receiving data from one or more cameras in an environment; generating metadata of one or more video analytics streams produced from the data; generating one or more time series of values based on the metadata; generating one or more affect scores for the one or more time series; generating a first signature representative of an observed behavior based on the one or more affect scores; performing pairwise matching by determining whether the first signature matches a second signature indicative of a query behavior; and performing an action when the first signature matches the second signature.

In another embodiment, a tangible, non-transitory computer-readable media is provided. The tangible, non-transitory computer-readable media stores computer instructions that, when executed by one or more processors, cause the one or more processors to: receive data from one or more cameras in an environment; generate metadata of one or more video analytics streams produced from the data; generate one or more time series of values based on the metadata; generate one or more affect scores for the one or more time series; generate a first signature representative of an observed behavior based on the one or more affect scores; perform pairwise matching by determining whether the first signature matches a second signature indicative of a query behavior; and provide an output when the first signature matches the second signature indicative of the query behavior.

In another embodiment, a system is provided that includes one or more cameras that capture data related to a behavior of one or more individuals in an environment. The system also includes one or more computing devices comprising one or more processors that receive the data from the one or more cameras; perform one shot learning using the data to determine whether an observed behavior of the one or more individuals matches a query behavior, wherein one shot learning comprises performing pairwise matching that is performed after only a single observation of an instance of the query behavior is obtained; and perform an action when the observed behavior matches the query behavior.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other features, aspects, and advantages of the present disclosure will become better understood when the following detailed description is read with reference to the accompanying drawings in which like characters represent like parts throughout the drawings, wherein:

FIG. **1** is an illustration of a social behavior recognition system, in accordance with an embodiment;

FIG. **2** is a flow diagram of a process suitable for generating a video analytics stream, in accordance with an embodiment;

FIG. **3** is a flow diagram of a process suitable for detecting behavior via one shot learning, in accordance with an embodiment;

FIG. **4** is a block diagram of example modules used to perform one shot learning, in accordance with an embodiment;

FIG. **5** is a matrix representation of thirteen behavior pairs used for one shot learning, in accordance with an embodiment;

FIG. **6** is a rank matrix based on thirteen observed behavior pairs and eighteen signal generators each capable of producing an affect score for each behavior, in accordance with an embodiment; and

FIG. **7** is a cumulative match characteristic curve for various one shot learning experiments, in accordance with an embodiment.

DETAILED DESCRIPTION

One or more specific embodiments will be described below. In an effort to provide a concise description of these embodiments, not all features of an actual implementation are described in the specification. It should be appreciated that in the development of any such actual implementation, as in any engineering or design project, numerous implementation-specific decisions must be made to achieve the developers' specific goals, such as compliance with system-related and business-related constraints, which may vary from one implementation to another. Moreover, it should be appreciated that such a development effort might be complex and time consuming, but would nevertheless be a routine undertaking of design, fabrication, and manufacture for those of ordinary skill having the benefit of this disclosure.

When introducing elements of various embodiments of the present disclosure, the articles "a," "an," and "the" are intended to mean that there are one or more of the elements. The terms "comprising," "including," and "having" are intended to be inclusive and mean that there may be additional elements other than the listed elements.

Embodiments of the present disclosure generally relate to a one shot learning framework for behavior recognition. One shot learning permits recognition of a behavior category or type on the basis of one or only a few behavior observations. Given a single observation of an instance of a query behavior (such as loitering), the social behavior recognition sys-